Grid Benchmarks and Workloads

Dror Feitelson

Hebrew University

Benchmarks

- An experiment designed to measure something specific
  - Microbenchmarks for isolated parameters
  - Full applications for overall performance
- A canonical workload
  - Agreed by the comunity
  - Used for comparisons when applied by different people
  - Linpack, SPEC, TPC, ...

Grid Workload

- Grid is a shared resource
- Performance on a loaded system may differ significantly from performance on a dedicated system
- Part of the infrastructure is responding to such variability (brokers, migration...)
- We need to be able to model and simulate these effects
- We need job-level grid workload data/model

Importance

- Lesson from MPPs:
  - Much effort spent optimizing single-job performance
  - Much lost to system-level fragmentation
- Needed to evaluate scheduling, resource management, co-allocation
- Need to do so in the context of competing local and grid jobs
  - Marks the difference between distributed and grid?

Workload Characterization

- Needs to be representative
  - Based on monitoring and logging
- Affected by local policies
  - Batch scheduling queues, runtime limits, hardware configuration
- Strange things happen in reality
- Often very modal
- Burstiness, self similarity, heavy tails
- Standard statistical techniques may be misleading

Take Home Message

- Please collect job-level data
- Please make it available
- Complement operational monitoring with the big picture context
- Understand system behavior even if not pathological
- Distinguish abnormal behavior from representative behavior

See Parallel Workloads Archive

http://www.cs.huji.ac.il/labs/parallel/workload/